The Case for the Relevancy of Downside Risk Measures

David Nawrocki
Villanova University
Department of Finance
College of Commerce and Finance
800 Lancaster Ave.
Villanova, PA 19085
610-489-7520 Voice or Fax
610-519-4323 Voice Mail
NawrockiD@aol.com **E-mail**

The Case for the Relevancy of Downside Risk Measures

I. Introduction

Why do we need downside risk measures like the semivariance and the lower partial moment in investment analysis? Very simply, we need these measures to cope with the complexity (and the reality) of the financial markets. The simple reason given in many articles on downside risk measures is that they are needed to deal with the skewness found in the nonnormal distributions of security returns.[1] This answer is too simplistic. We need downside risk measures because they are a closer match to how investors actually behave in investment situations. The theory of economic behavior is known as utility theory. It states that economic units will act to maximize their economic satisfaction (or utility). Utility theory has rarely been taught in finance courses the past thirty years because the market theories developed in the 1960s effectively eliminated the need for economic utility theory. The main culprit is the Capital Asset Pricing Model which makes beautiful intuitive sense but, unfortunately, has no grounding in the reality of how financial markets actually work. The past thirty years has provided no empirical support for the CAPM. In fact, the academic finance profession has known for twenty years that it was not feasible to test the CAPM.[2]

Imagine Monty Python's famous "parrot sketch" as a conversation between a finance academic and a finance practitioner. John Cleese plays the part of a finance practitioner and Michael Palin plays the part of a finance professor.

The CAPM Sketch

John Cleese (Entering a large university): "Excuse me. I would like to register a complaint about this financial market theory which I purchased from this very boutique."

Michael Palin: "Ah, the CAPM, a remarkable theory. What's wrong with it?"

John Cleese: "I will tell you what is wrong with it. It's dead."

Michael Palin: "Nah, Nah. It's resting. It will be up and about shortly. Haven't you been reading "The Journal"?

John Cleese: "Never mind that. I know a dead theory when I see one and I'm looking at one right now."

Michael Palin: "Nah, it's not dead. It's resting."

John Cleese (Incredulous): "Resting?"

Michael Palin: "Yeah, resting. A remarkable theory, the CAPM. Beautiful plumage."

John Cleese (Raising his voice): "Beautiful plumage? It's stone dead."

Michael Palin: "No, it's just resting."

John Cleese: "This theory is definitely deceased. This theory wouldn't move if you put 4000 volts through it. It's bleeding demised. It's passed on. This theory is no more. It has ceased to be. It has expired and gone to meet its maker. This is a late theory. It's bereft of life. It rests in peace. If you weren't still trying to salvage it, it would be pushing up the daisies. It has rung down the curtain and joined the choir invisible. This is an ex-theory."

Michael Palin: "Well, then we'll have to replace it." (After rustling around the university for a few moments, he returns) "I'm sorry but we don't have any more financial market theories left."

Let's leave John Cleese and Michael Palin and ponder the question: "Why is the CAPM still taught in finance textbooks and CFP courses?" The answer is "It reduces the complexity of the markets down to a few simple rules." All encompassing theories like the CAPM make the markets easier to understand. The problem is that the CAPM model does not mirror reality. The price of simplicity is that the model is not relevant to the real world. If the financial markets operated according to CAPM, we would simply purchase the market index portfolio and manage risk and return by mixing risky stock portfolios with riskfree treasury bills.

Other than the simple rules derived from CAPM, we wouldn't have to have any other substantial knowledge. We wouldn't have to study law, accounting, tax code, macroeconomics, business cycles, human behavior, human decision-making, utility theory, psychology, social psychology, socialization, philosophy, product quality, marketing, distribution channels, design, management, etc. Would a good investment advisor buy a client a market index mutual fund in November that is up 35% year to date and will be distributing a ton of short term capital gains to shareholders in the near future? The answer: "It depends on the tax status of the client". Did Gibson Greetings and Procter and Gamble understand the legal contracts (known as derivatives) that Bankers Trust was selling to them? In a CAPM financial world, none of this knowledge is relevant. The market magically takes care of these mundane details.

In a CAPM market, how does a financial professional add value in order to justify the fees or commissions charged to customers? It can't be done. The financial professional is basically a ticket office charging admission to the financial markets essentially similar to the ticket office at the movie complex. If the market is not efficient and if it does not operate according to CAPM, then the financial advisor can add value by providing knowledge. The financial advisor has to provide substance, i.e., something to bring to the table. That substance is knowledge: knowledge of law, accounting, taxes, economics, human behavior, etc.[3]

The little piece of substance in which this paper is interested is the application of utility theory through the use of downside risk measures.

There are two serious problems with using CAPM to build asset allocation models.

First is the diversification problem. Unfortunately, the investment cannot be in one security as this exposes the investor to default or bankruptcy risk and business risk. The response to this problem is to diversify the investment into many different investments, which will lower the overall return. Diversification using a large number of assets over a long period of time is a very complex multi-dimensional dynamic programming problem. The probability that an investor can solve this problem without a computer is zero. The problem is the multi-dimensionality. If the portfolio is going to be built from 100 individual assets, the programming problem will involve 100 dimensions plus additional dimensions to handle the time horizon in the dynamic programming problem.

This results in a very serious problem for the investor. Very simply, humans do not think well beyond three dimensions. Our minds are limited to the three dimensions that we can sense. Therefore, we will have a problem allocating funds to 5 assets over a 5-year time horizon. The computer does not have this limitation and can mathematically solve a complex multi-dimensional problem.

A second problem is the ability to instill a complete picture of the investor's goals, aspirations, expectations, etc. into one utility function that can be solved by our multi-dimensional dynamic programming. The behavior of a human cannot be distilled into one utility function, but rather a multitude of utility functions is required to describe the behavior of an individual.

Still, investors do make allocation decisions without resorting to 100 plus dimensional computer programs and aggregate utility functions and the solutions seem to work. Why? The major barrier to an understanding of investor behavior is the concentration of attention on the behavior of an idealized investor in a highly constrained environment, i.e. the perfectly rational investor. This results in a model of ***how investors are supposed to behave*** given numerous simplifying assumptions so that the rational investor can maximize returns and minimize risk. If every investor in marketplace behaves according to rational investor return maximization and liquidity risk minimization, it would wreak havoc in the marketplace. Weiner (1948) states that the market would be highly volatile careening from overbuying to overselling. The aggregate behavior of rational investors will create a monster roller coaster ride for the markets. Weiner (1948) suggests that humans will group together into cooperatives (savings banks, credit unions, savings and loan, mutual insurance funds, and mutual funds) to reduce the uncertainty of the financial marketplace and to protect the group from the rational investors. Peters (1992) describes a market dominated by short-term time horizon investors as highly unstable with huge volatility. Both Weiner (1948) and Bear and Maldonado-Bear (1994) state that society will have to pass laws to protect society in general from the behavior of these rational investors. Of course, the USA has passed numerous security laws to protect society from rational behavior. If rational behavior is so unacceptable to society in general that it legislates against this behavior, how realistic is it to assume all investors in the marketplace are rational?

The key to understanding the operation of financial markets is to understand ***how investors actually do behave*** in the financial marketplace. Since an aggregate (one) utility function is impossible to derive for a human investor, it is pretty certain that investors do not use aggregate utility functions. The alternative to aggregate utility functions was being developed in the 1950s and 1960s when the financial market theories such as CAPM came along and swept away everything in their paths.

It is impossible that one all-encompassing aggregate utility function that will work for a person's entire life can be derived. As a person lives their life, their goals, tax situation, etc. are constantly changing, therefore, the person's utility function is always a work in progress.

The first author to address this problem was Simon's (1954) work on utility satisficing. Simon states that humans will not optimize their utility but rather will accept satisfactory results from a limited search rather than an optimal search. Cyert and March (1963) followed with probably one the best books on corporate finance, *The Behavioral Theory of the Firm*. Cyert and March studied human behavior. From these studies, they generated a model of how humans within organizations make decisions. They state that complex problems are solved by breaking down multi-dimensional problems into a series of two-dimensional problems and then solving these problems sequentially until a satisfactory solution is achieved.

We also can proceed by breaking the problem into subproblems and achieving satisfactory results with each subproblem. This procedure is going to be controversial with most academics who believe in the possibility of the aggregate utility function, but the practitioner has to be pragmatic and use techniques that are possible.[4] One problem cited by academics is that the process of solving subproblems over time will be myopic, i.e. shortsighted. However, this argument is based on an assumption that the person and the market environment is stationary and never experiences change. In reality, with complex changing environments, the investor has to take Weiner's (1948) advice and engage in adaptive behavior based on adaptive feedback controls.

This procedure is reflected in the compartmentalization of utility. An individual is going to have different financial compartments, each with a different goal, a different utility function, and a different solution. When we aggregate the results of the compartmentalization process, we will achieve a satisfactory result not an optimal result.

The individual's financial situation is broken into compartments. Each compartment has a different goal and time horizon. Each compartment has a different utility function. Each compartment will have a different solution. As a goal is achieved, this affects the remaining goals. Therefore, the compartmentalization process has to be repeated on remaining goals. As a person moves through their life, their financial situation changes and their allocation decision will have to be continually re-solved.

As a person changes over time, the investment environment is also changing. As a result of a changing human being and a changing environment, there is no such thing as a static utility function or a static financial plan. Revising and re-solving the compartment problem can help an individual keep up with their changing conditions.

The environment changes because economic units go through a life cycle. Businesses go through a product life cycle. As a result, firms will go through a life cycle from being a startup firm, to a high growth firm, to a cash cow and, finally, to a profit challenged firm heading to liquidation.

Measuring risk under these conditions of investor compartmentalization of investment goals and business cycles is not going to be that difficult. First, statistical measures of risk are going to work because the constant revision and re-solving the compartment problem can lead to relatively short-term time horizons. Using statistics like the standard deviation and the semivariance measures the liquidity risk of an investment and is only relevant to a short-term investment horizon investor. Does a standard deviation calculated from 20 years of data provide any useful information? Looking backwards and taking a historic perspective, the standard deviation provides a meaningful interpretation of history. However, looking forward for the next one, two, five, or ten years, the information in the standard deviation is basically useless because it only measures short-term liquidity risk. If the investor has a long-term investment horizon (say 20 years), then it is safe to ignore liquidity risk and simply maximize the expected return of the investment.

Second, the risk measure has to reflect the investor's utility, the allocation of funds in the investor's portfolio has to reflect the investor's utility, and the measurement of the investment performance has to reflect the investor's utility. For all of these steps, a risk measure that approximates liquidity risk and investor utility can be used. The first candidates for the appropriate risk measure are the beta and the standard deviation. Both, unfortunately, represent only one utility function which provides a "one size fits all" approach. The second candidate is the Lower Partial Moment (LPM) risk measure which provides a multitude of utility functions that represent the whole range of human behavior, i.e. risk seeking to risk neutral to risk averse.

The LPM is computed using different degrees. The degree, n, represents the investor's utility in terms of risk aversion. When $n<1$, the investor is a risk seeker. When $n=1$, the investor is risk neutral. When $n>1$, the investor is averse to risk. The higher the value of n, the higher the level of risk aversion. Within the utility theory literature, individuals with degrees of risk aversion as high as 4.0 have been identified.[5]

The n-degree LPM and the compartmentalization of utility concept can be used to explain complex, seemingly contradictory, human behavior. An example is a male, who is approaching retirement age and whose total lifetime savings is $500,000. Our hero has an appointment to see a financial advisor and at the meeting agrees to place the $500,000 in a bond portfolio consisting of AAA rated corporate bonds and US Treasury bonds.

Since the $500,000 represents all of his wealth, the portfolio represents a very high degree of risk aversion. He signs the papers to implement this plan. As our hero leaves the office building where he met with the financial advisor, he stops into a store and buys a daily number lottery ticket for $1.00. Compartmentalization explains that this behavior is not irrational, but rather each decision is rational within its compartment. In the first compartment, the investor is very risk averse since all of his wealth is at stake. The lottery ticket in the second compartment represents no threat to the investor's wealth as long as it is only $1.00 every so often. (When a person is buying $20-30 worth of lottery tickets daily, 5 or 6 days a week, then this is probably a matter for Gambler's Anonymous. Again, this depends on the person's total wealth or income.) The investor can be risk seeking in the second compartment because the wealth amount is so low relative to his total wealth. Any entertainment such as movies, theatre, sporting events, horse racing, casinos, amusement parks, etc. could be included in this category. The lottery ticket, in this case, is entertainment not an investment.

At this point, a financial planner has every right to be confused. For the most of the past twenty years, the academics have been attempting to have financial planners derive an aggregate solution for clients based on an aggregate utility function. A very extensive asset class optimization solution (asset allocation) has been recommended as the proper implementation of portfolio optimization programs. The typical mean-variance optimization of asset classes assumes that all clients have a short-term quadratic utility function. By changing the slope of the quadratic utility function, different risk-return tradeoffs may be chosen for the client. However, the solution is for short-term risk aversion and represents one aggregate utility function for the client. Figure 1 represents the utility behavior of three investors. Each investor has a quadratic utility function, the only difference being that their utility curves have different slopes indicating different levels of risk aversion. The slope is the *change in the Y-axis (Return) divided by the change in the X-axis (Risk).* Investor A is willing to accept larger amounts of risk in exchange for small increases in return. Investor B is balancing risk and return somewhat equally. Investor C is willing to give up larger returns in exchange for smaller amounts of risk. In all of these cases, risk is expressed through the variance measure.

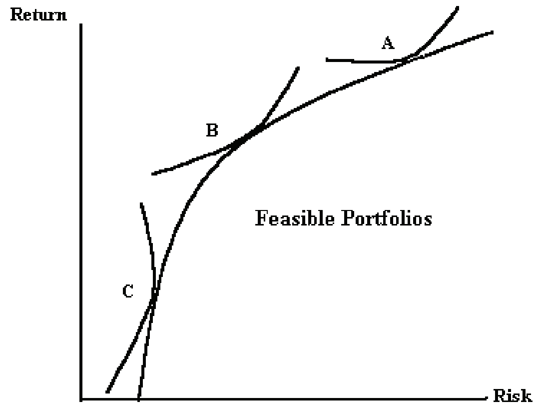

Figure 1 -- Traditional Risk-Return Analysis with Quadratic Utility Functions

The short-term risk aversion inherent in an efficient frontier analysis is a problem. Actually, the risk-return feasible space is changing over time. Think of the market over time as a stick of pepperoni. Each time a pepperoni slice is sliced off to put on a pizza, the cross section of the slice is a point in time. The markets can be described as a series of pepperoni slices (or cross sections) over time as in Figure 2.

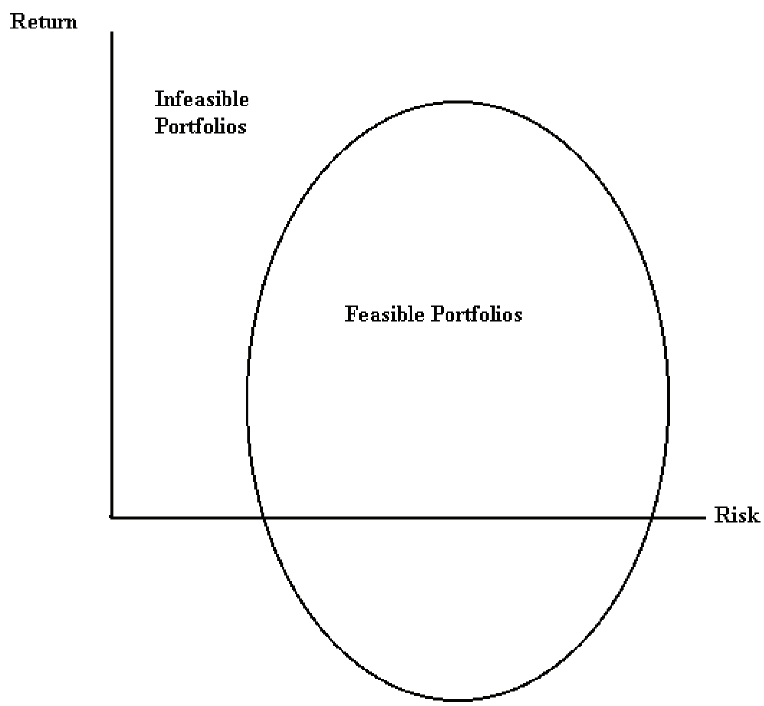

Figure 2 -- Risk-Return Space During One Moment in Time

Therefore, the traditional analysis is only looking at short-term liquidity risk at one point in time. The variance risk measure cannot handle long term exposure to risk except through the rather heroic assumptions of stationary expected returns and variances over the time period. As this is not realistic, it is not a very helpful solution. Asset allocation packages, however, still provide an analysis called time diversification where these heroic assumptions are used to develop a long-term investment plan for the client.

It is imperative that the use of asset allocation not be confused and misused. There are two ways to use asset allocation appropriately. First, the client may segment their funds by different goals. Second, a client may segment investments into different segments of the markets depending on dynamic nature of the market.

Asset allocation may be used with different asset classes to provide a diversified portfolio to meet client needs for a single investment need. Asset allocation should not be used to try to meet multiple goals from different investor utility compartments. The compartmentalization of utility approach allows each compartment to have a different time horizon, a different utility function, and a unique solution. The compartments are assumed to be independent of each other, however, interrelationships between some of the compartments can be handled within the Cyert and March sequential solution of subproblems framework.

A framework is also needed to handle the dynamic nature of the market environment and help to improve planning. One framework that seems promising is the product life cycle, which is used by many security analysts to study rapidly changing industries. A firm can go through many stages of a product life cycle depending on its product mix. If a firm is a single product *startup* or has many new products, it is early in its product life cycle. Next, a successful product will go through a strong *growth* period. Early in the growth period there will be large growth rates and low or negative profits. Late in the growth period, the growth rate will start to decline but the profits will increase. During the *cash cow* period, the firm experiences growth rates consistent with the general economy's growth rate. However, the firm has a mature product market and generates large profits and high cash flows. Finally, the product profitability turns to losses and if new products do not come along to restart the firm's product life cycle, the firm will move towards *liquidation.*

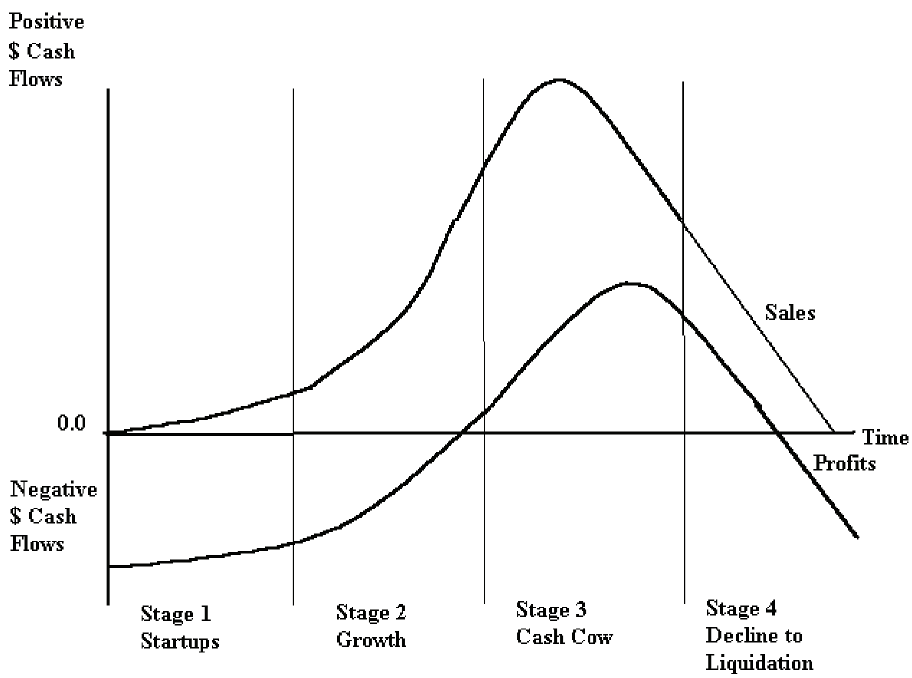

Figure 3 -- The Four Stages of a Corporation's Product Life Cycle

The product life cycle (Figure 3) is where the utility functions in the n-degree LPM become important. In the early *Startup* stage, investors will have to exhibit high degrees of risk tolerance and may actually engage in risk loving behavior, i.e. taking on high degree of risk for the small chance for a large return. While the short-term risk is very high and current returns are negative, these investors are interested in startups because of their long term potential. When companies start their *Growth* stage, returns are still negative and risk is still high but the odds of the firm experiencing strong growth to positive cash flows is improving. Again, investors will be willing to exchange short-term risk and negative returns for potentially high future returns. Later, in the *Growth* stage, firms will be profitable and more investors will be attracted to the company because of high returns and lower risk. The *Cash Cow* stage is where the company has established its franchise within the marketplace and is generating strong cash flows. *Cash Cows* have typically achieved full market penetration and will experience low growth rates tied to the general population growth rate. Finally, the company's franchise will wear down, profitability will drop, and speculators will start hovering around the company betting on the liquidation value of the company. Now the firm has entered the last phase of its life cycle, the *Liquidation* phase. If the firm can successfully develop new products to start its life cycle over again, the firm can remain in business. Otherwise, it will be liquidated.

Figure 4 demonstrates the use of the n-degree LPM measure. The degree n is a measure of the investor's attitude towards risk. When the efficient frontiers derived from different degrees of the LPM are plotted on one graph, the problem of measuring risk on the X-axis arises. In both graphs, the LPM n=2 measure is used on the X-axis. Therefore, the efficient frontier derived using the LPM n=2 will be the dominant or best frontier. The more risk averse frontiers (n=3 and n=4) will be less efficient subsets of the n=2 frontier. The risk neutral (n=1) and risk loving (n=.6 and n=.2) will also be less efficient subsets of the n=2 frontier. Note that the risk loving frontiers will experience the largest increases in risk as measured by the LPM n=2. Now if we were to use LPM n=4 on the X-axis, the n=4 frontier will be the dominant frontier and the other frontiers will be less efficient subsets of the n=4 frontier. As we change the risk metric on the X-axis, the dominant frontiers will always be the frontier derived using the same risk measure as the risk measure on the X-axis.

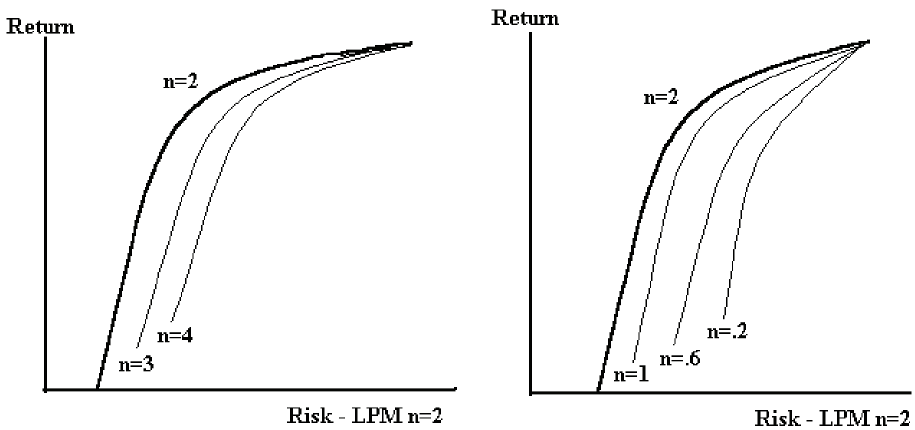

Figure 4 -- Efficient Frontiers Generated by LPM Optimizer Using Different Values of n-Degree LPM. All Efficient Frontiers are Graphed Using LPM n=2 as the Risk Measure.

The previous discussion is necessary to understand Figure 5. First, the risk measure on the X-axis is the LPM n=1 risk neutral measure. Next, the feasible frontier is segmented using the product life cycle. *High risk and negative returns will typically characterize startups*. *Growth* companies will run the range from high risk and negative returns to high risk and positive returns. *Cash Cows* will be represented in the low risk-positive return section of the graph. *Liquidation* candidates will be moving towards the higher risk-negative return area of the graph.

With *Startup* companies, the appropriate utility function will be a risk seeking function (n<1.0). Investors in startup companies typically will invest in a large number of startups over a long period of time hoping to hit one home run. Anyone looking at a startup on a short time horizon risk-return graph is going to see a very risky investment with negative returns, therefore, a risk seeking utility function would be used. Although, the investor in a startup is usually looking at a long-term investment horizon, the short-term risk seeking utility function will be used in the risk-return analysis to correct for the short-term nature of a statistical risk measure. If we limit the portfolio optimization to only *Startup* companies, the efficient frontier (**AB**) for LPM n=1 will be the dominant frontier because the X-axis is measured in units of LPM n=1. Risk loving frontiers (n=0.8, n=0.6, and n=0.2) will be subsets of the **AB** n=1 frontier. However, for an investor with a risk loving utility function of n=0.8, the n=0.8 frontier will be the efficient frontier. From there, the investor will have to pick one portfolio that matches the investor's risk-return profile.

Investors in high *Growth* companies are looking at firms with high risk and variable returns. Firms early in their *Growth* phase will have negative returns while firms later in the growth phase will have higher returns. Investors will have long-term investment horizons and will be risk loving (n<1.0) and will invest in the *Early Growth* frontier, **CD**, and its subsets (n=0.8 is one example). Investors who are risk neutral and slightly risk averse will invest in the *Late Growth* frontier, **EF,** and its subsets (n=1.2 and n=1.4)

When firms are experiencing high growth rates, it's because they are continuing to open new markets to their product. As the industry grows, it will reach saturation points where all markets have access to the product and the market is saturated with the product. At this point, high growth rates cannot continue and will slow down to the general market growth rate. At this point, the firm becomes a *Cash Cow.*

Investors purchasing *Cash Cows* are interested in current income and liquid investments. Their utility function will be short term and will be risk averse (n>1.0). They will invest in the *Cash Cow* frontier, **GH**, and its risk averse subsets (n=2, n=3, and n=4).

Finally, firms headed for *Liquidation* will attract investors with long term investment horizons who are betting that the liquidation value of the firm is more than the current market value. Again, these investors will have short-term risk seeking utility functions (n<1.0) and will invest in the *Liquidation* frontier, **IJ**, and its risk seeking subsets (n=0.8 and n=0.6).

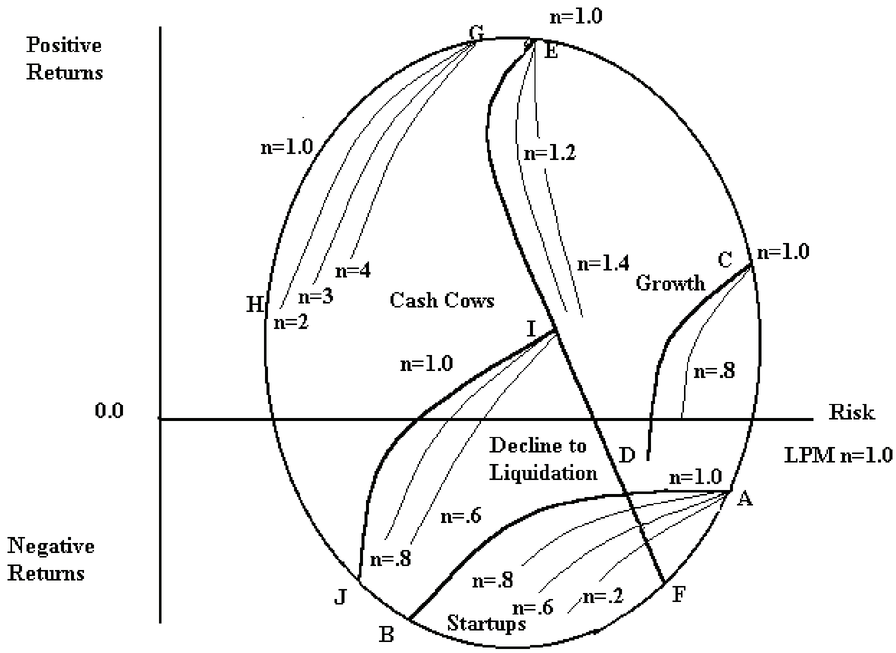

Figure 5 -- Location of Firms in Different Stages of Product Lifestyles in a Single Period Risk-Return Space

The n-degree LPM allows the investment manager to compartmentalize utility. For example, a portfolio algorithm can be used to solve for a portfolio of startup companies for a value of $n<0$. A portfolio algorithm can be used to solve for a portfolio of growth companies for a value of $n=1$. A portfolio algorithm can be used to solve for a portfolio of cash cow companies for some value $n>2$. There will be a unique efficient frontier for each individual degree, n, of the LPM measure. Therefore, there is a unique LPM efficient frontier for each compartment of an individual's investment decision.

It is this inherent ability of the n-degree LPM to fit into the process of how investors actually make investment decisions that makes the LPM downside risk measures so important to the investment community. It describes ***how investors actually do behave*** rather than ***how investors are supposed to behave***.

Footnotes

1. It would be nice to write an article without a single equation. The LPM measures the downside risk below some target return set by the investor. The equations for the variance and the LPM are as follows:

$$V_i = \frac{1}{m} \sum_{t=1}^{m} (R_t - E(R))^2 \qquad (1)$$

$$LPM_{in} = \frac{1}{m} \sum_{t=1}^{m} [\text{Max}(0, h - R_t)]^n \qquad (2)$$

   Where V is the variance for security **i** with **t** = 1, 2, ...**m** observations. R is the return for period **t** and E(R) is the expected mean return for security **i**. The square root of the variance is the standard deviation. The LPM is the lower partial moment for security **i** and for degree **n** with **t** = 1,2,......**m** observations. Max is the maximization function that selects the larger of two numbers: either **0** or $\mathbf{h-R_t}$ and **h** is the target return for the portfolio. The degree **n** determines the power exponent of the differences. When **n** = 2, then the lower partial moment is known as the semivariance. Taking the square root of the LPM n=2 will provide a downside risk measure known as the semideviation. The LPM can also be raised to the **n** power and the **n**th root can be taken. A number of the traditional statistical advantages of the LPM measure can be found in Balzer (1994), Rom and Ferguson (1993) and in Sortino and Van Der Meer (1991). Merriken (1994) emphasizes that risk measures such as variance and LPM are appropriate for investors with the short-term investment horizons.

2. Roll (1977) wrote one of the first critiques of the CAPM model. His final conclusion was that the CAPM model was not testable, therefore, there is no empirical support for the model.

3. This section derives from the discussion in Frankfurter (1997).

4. This view, however controversial, is supported by academic research. There is a strong academic body of theory known as evolutionary economics that has integrated systems theory and economics. Boulding (1981) was the major proponent of this school of thought. The integration of systems theory and the Cyert and March (1963) behavioral theory into evolutionary economics is presented in Nelson and Winter (1982).

5. Fishburn (1977) provides the theoretical support for using lower partial moment to capture the individual utility function of a specific investor.

References

Balzer, Leslie A. (1994). "Measuring Investment Risk: A Review," Journal of Investing, v3(3), 47-58.

Bear, L.A. and R. Maldonado-Bear (1994). Free Markets, Finance, Ethics, and Law, Englewood Cliffs, NJ: Prentice Hall.

Boulding, K.E. (1981). Evolutionary Economics, Beverly Hills, CA: Sage Publications.

Cyert, Richard M. and James G. March (1963). A Behavioral Theory of the Firm, Englewood Cliffs, NJ: Prentice-Hall.

Fishburn, Peter C. (1977). "Mean-Risk Analysis With Risk Associated With Below-Target Returns," American Economic Review, v67 (2), 116-126.

Frankfurter, George M., ed. (1997). "Pushing the Epsilon to the Abyss Post-Modern Finance," Panel Discussion, International Review of Financial Analysis, v6, (2), 133-177.

Merriken, Harry E. (1994). "Analytical Approaches To Limit Downside Risk: Semivariance And The Need For Liquidity," Journal of Investing, v3(3), 65-72.

Nelson, R.R. and S.G. Winter (1982). An Evolutionary Theory of Economic Change, Cambridge, MA: Harvard University Press.

Peters, E.E. (1991). Chaos and Order in the Capital Markets, New York: John Wiley and Sons.

Roll, Richard (1977). "A Critique Of The Asset Pricing Theory's Tests; Part I: On Past And Potential Testability Of Theory," Journal of Financial Economics, v4, (2), 129-176.

Rom, Brian M. and Kathleen W. Ferguson (1993). "Post-Modern Portfolio Theory Comes Of Age," Journal of Investing, Winter 1993, Reprinted Fall 1994, v3 (3), 11-17.

Simon, Herbert A. (1955). "A Behavioral Model Of Rational Choice," Quarterly Journal of Economics, v69 (1), 99-118.

Sortino, Frank A. and Robert Van Der Meer (1991). "Downside Risk," Journal of Portfolio Management, v17 (4), 27-32.

Weiner, N. (1948). Cybernetics, Cambridge, MA: The MIT Press.